A tight distance-dependent estimator for screening three-center Coulomb integrals over Gaussian basis functions

David S. Hollman,[1,2] Henry F. Schaefer,[1] and Edward F. Valeev[2]
[1)] *Center for Computational Quantum Chemistry, University of Georgia, 1004 Cedar St., Athens, Georgia 30602, USA*
[2)] *Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA*

A new estimator for three-center two-particle Coulomb integrals is presented. Our estimator is exact for some classes of integrals and is much more efficient than the standard Schwartz counterpart due to the proper account of distance decay. Although it is not a rigorous upper bound, the maximum degree of underestimation can be controlled by two adjustable parameters. We also give numerical evidence of the excellent tightness of the estimator. The use of the estimator will lead to increased efficiency in reduced-scaling one- and many-body electronic structure theories.

arXiv:1409.8232v1 [physics.chem-ph] 29 Sep 2014

I. INTRODUCTION

Efficient evaluation of matrix elements of the Coulomb operator (electron repulsion integrals, ERIs) is an essential ingredient of practical quantum chemistry, especially in the context of one-body methods [Hartree–Fock (HF)[1] and hybrid Kohn-Sham density functional theory (DFT)[2]] and lower-order many-body methods [second-order Møller–Plesset (MP2)[3]]. The ERI tensor is defined as

$$g_{AB} \equiv (A|B) = \int \Omega_A(\vec{r}_1)\frac{1}{r_{12}}\Omega_B(\vec{r}_2)d\tau_1 d\tau_2, \quad (1)$$

where Ω_A and Ω_B are the (Mulliken notation) bra and ket charge distributions. The most common version of the ERI tensor is the four-center case, in which both Ω_A and Ω_B are product densities, composed of either atom-centered basis functions in HF/DFT or molecular orbitals in correlated methods such as MP2. To reduce the computational cost of the four-center ERIs, Ω_A and Ω_B can be approximately expanded in terms of auxiliary basis set; the resulting density-fitting (DF) (also known as the resolution-of-the-identity, RI) approximation involves also three- and two-center ERIs in which either or both of the bra and ket densities are single, atom-centered functions.

When spatially localized basis functions are used, the number of ERIs with magnitude greater than some threshold ϵ is $\mathcal{O}(N^2)$ rather than $\mathcal{O}(N^4)$, where N is the size of the basis set (and thus proportional to the size of the molecule). To take advantage of this sparsity it is necessary to estimate the magnitude of ERI very rapidly (i.e., significantly faster than it would take to compute it rigorously). The simplest such estimate invokes the Schwarz inequality[4–8]

$$|(A|B)| \leq |(A|A)|^{\frac{1}{2}} |(B|B)|^{\frac{1}{2}}, \quad (2)$$

which holds because Eq. (1) defines an inner product.

While the Schwarz bound substantially reduces the cost of computing the ERI tensor by screening out small Ω_A and Ω_B, it is not a tight bound because it neglects the decay over distance between Ω_A and Ω_B. Extensive research has been done recently on integral estimates that incorporate the distance between the bra $(A|$ and the ket $|B)$ distributions.[9–13] Earlier approaches were based on the multipole expansion of the Coulomb kernel,[9–11] but these efforts were later abandoned in favor of an empirical modification of the Schwarz estimate, referred to as the QQR estimator.[12,13] For four-center ERIs, the QQR estimator can be written as

$$|(\mu\nu|\lambda\sigma)| \approx \frac{|(\mu\nu|\mu\nu)|^{\frac{1}{2}} |(\lambda\sigma|\lambda\sigma)|^{\frac{1}{2}}}{R - \mathrm{ext}_{\mu\nu} - \mathrm{ext}_{\lambda\sigma}} \quad (3)$$

where $\mathrm{ext}_{\mu\nu}$ and $\mathrm{ext}_{\lambda\sigma}$ are the well-separatedness (WS) extents from the continuous fast multipole method (CFMM),[14] and R is the distance between centers of charge of bra and ket.

Although the QQR estimator is not an upper bound, it is a tight estimate. In practice, the tightness is more important than the upper bound property anyway because the precision of the computed property (such as energy) is a complex function of the truncation threshold for the ERIs. Whereas some have investigated how to compute, for example, the Fock matrix with guaranteed precision in a given Gaussian basis,[15,16] guaranteeing precision of the energy and other properties is even more difficult due to their nonlinear dependence on the Hamiltonian via the density matrix/wave function, among other factors. In practice the relationship between precision of the target property and the ERI truncation threshold is established empirically. Underestimation by the ERI estimator will cause some integrals that normally would have been deemed important enough (above the threshold) to be skipped. However, this will only affect the small integrals near the threshold. Furthermore this can be accounted for by adjusting the empirical relationship between precision and truncation threshold. Thus, the lack of the upper-bound property is not an issue in practice. More discussion of the advantages and disadvantages of rigorous upper bounds versus reliably tight bounds can be found in Ref. 12.

In the context of the MP2 method it is necessary to replace the "effective" distance in the denominator of the QQR estimate with its second or third power to take into account the vanishing leading-order multipoles of

bra and ket densities.[13] The same issue arises when we want to estimate the magnitude of the three-center integrals in the context of the context of reduced-scaling electronic structure methods that utilize (local) density fitting approximations.[17–32] Indeed, Eq. (3) is not a good estimator for the two- and three-electron integrals over atom-centered Gaussian basis functions because one or more of the lowest-order multipoles of Ω_A and/or Ω_B will vanish for many integrals, and hence long-range decay with R will be significantly faster.

To take advantage of the sparsity of the three-center integrals, we developed a rapid estimator with proper asymptotic R-dependence. Although our estimator differs technically from the QQR estimator of Ochsenfeld and co-workers,[12] it is similar in spirit. The primary objective of the estimator is tightness, not a rigorous upper bound. However, the underestimates in our new bound can be reduced arbitrarily by adjusting the WS parameter. The estimator is also governed by a second parameter, ϑ_{SQ}, which loosely controls the trade-off between tightness of the estimator and the maximum amount of underestimation.

II. THEORY

In this work, we will deal with the basis of atom-centered contracted real solid harmonic Gaussian-type orbitals (GTOs), defined as[33,34]

$$\chi_\mu(\vec{r};\vec{R}) = \sum_a c_{a,\mu}\phi_{a\ell_\mu m_\mu}(\vec{r}-\vec{R}), \tag{4}$$

$$\phi_a(\vec{r}) = \mathcal{S}_{\ell_a m_a}(\vec{r})\exp(-\zeta_a|\vec{r}|^2) \tag{5}$$

$$\mathcal{S}_{\ell m}(r,\vartheta,\varphi) = (-1)^m N_{\ell m}(\zeta_a) \times r^\ell P_\ell^{|m|}(\cos\vartheta)\begin{cases}\cos(m\varphi) & m\geq 0\\ \sin(|m|\varphi) & m<0\end{cases} \tag{6}$$

$$N_{\ell m}(\zeta_a) = \left[\frac{\pi\Gamma(\ell+\frac{3}{2})(\ell+|m|)!}{(2\zeta_a)^{\ell+\frac{3}{2}}(2\ell+1)(\ell-|m|)!}(1+\delta_{m0})\right]^{-\frac{1}{2}} \tag{7}$$

where $c_{a,\mu}$ are the contraction coefficients, P_ℓ^m are the associated Legendre polynomials, and $\Gamma(\ell+\frac{3}{2})$ is the gamma function [since ℓ is a nonnegative integer, $\Gamma(\ell+\frac{3}{2}) = \sqrt{\pi}(2\ell+1)!!/2^{\ell+1}$].

A. Notation

Throughout this work we will use Mulliken (chemists') bra-ket notation shown in Eq. (1). Three-center ERI will be written with the principal (orbital) basis product $(\mu\nu|$ in the bra and the one-center auxiliary (density fitting) basis function $|X)$ in the ket. Greek letters μ and ν will be used to denote functions in the principal basis, and X will be used to denote basis functions in the auxiliary basis. For primitives, we will use a and b for the principal basis and c for the auxiliary basis.

B. The Primitive $(ss|X)$ Case

Consider a three-center ERI $(ab|c)$, where primitive GTOs ϕ_a, ϕ_b, and ϕ_c have principal angular momentum quantum numbers $\ell_a = 0$, $\ell_b = 0$, and arbitrary ℓ_c. We use the notation m_a, m_b, and m_c for the L_z quantum numbers; ζ_a, ζ_b, and ζ_c for the primitive exponents; and $\vec{R}_a$, $\vec{R}_b$, and $\vec{R}_c$ for the centers of the respective primitives (analogous notation will be used throughout). According to the Gaussian product theorem (GPT), the bra $(ab|$ can be rewritten as a single Gaussian centered at the center-of-exponent point $\vec{R}_p$ located on the line in between $\vec{R}_a$ and $\vec{R}_b$ (temporarily neglecting normalization):

$$\phi_a(\vec{r}_a)\phi_b(\vec{r}_b) = \exp(-\zeta_a|\vec{r}_a|^2)\exp(-\zeta_b|\vec{r}_b|^2) \tag{8}$$

$$= K_{ab}\exp(-\zeta_p|\vec{r}_p|^2) \tag{9}$$

$$\equiv \Omega_A(\vec{r}_p), \tag{10}$$

where $\zeta_p = \zeta_a+\zeta_b$, $\vec{R}_p = (\zeta_a\vec{R}_a+\zeta_b\vec{R}_b)/\zeta_p$, $\vec{r}_i \equiv \vec{r}-\vec{R}_i$, and $K_{ab} \equiv e^{-\zeta_a\zeta_b|\vec{R}_a-\vec{R}_b|^2/\zeta_p}$. Since Ω_A is spherically symmetric, from a large enough distance away it can be viewed as a point charge centered at $\vec{R}_p$ of magnitude S_{ab} (overlap between ϕ_a and ϕ_b).

Now consider the potential generated at $\vec{R}_p$ by a primitive Gaussian ϕ_c located at $\vec{R}_c$. For practical purposes, we wish to determine the average potential of an entire shell (i.e., all m_c with $-\ell_c \leq m_c \leq \ell_c$) rather than an individual m_c. Thus, without loss of generality we can orient our system with the z axis along the vector $\vec{R} \equiv \vec{R}_p - \vec{R}_c$. With this orientation, a nonzero potential is generated only by the $m_c = 0$ Gaussian, and we are left with (omitting normalization of ϕ_c for now)

$$\tilde{V}_{\ell_c}(R) = \int \frac{r^\ell}{|\vec{r}-\vec{R}|}\exp(-\zeta_c r^2)P_{\ell_c}(\cos\vartheta)d\vec{r}. \tag{11}$$

We then use the multipole (Laplace) expansion of the Coulomb operator

$$\frac{1}{|\vec{r}-\vec{R}|} = \sum_{k=0}^{\infty}\frac{r_<^k}{r_>^{k+1}}P_k(\cos\vartheta), \tag{12}$$

where ϑ is the angle between $\vec{r}$ and $\vec{R}$, $r_< = \min(R,r)$,

and $r_> = \max(R, r)$, to obtain the potential:

$$\tilde{V}_{\ell_c}(R) = \frac{4\pi}{2\ell+1}\left[\int_0^R \frac{r^{2\ell+2}}{R^{\ell+1}}\exp(-\zeta_c r^2)dr + \int_R^\infty R^\ell \exp(-\zeta_c r^2)dr\right] \tag{13}$$

$$= \frac{4\pi}{2\ell+1}\left[\frac{\Gamma\left(\ell_c+\frac{3}{2}\right)-\Gamma\left(\ell_c+\frac{3}{2},R^2\zeta_c\right)}{2R^{\ell+1}\zeta^{\ell+\frac{3}{2}}} + \frac{\exp(-\zeta_c R^2)R^{\ell_c}}{2\zeta_c}\right], \tag{14}$$

where $\Gamma(s,x)$ is the (upper) incomplete gamma function. The second term in Eq. (14) rapidly decays to zero for reasonable values of ζ_c. Since we are interested in behavior at large R (i.e., where the bra can be viewed as a point charge), we can omit this term. Similarly, the incomplete gamma function $\Gamma(s,x)$ is bounded from above by $\Gamma(s)$ for real, positive x, and the second part of the first term in Eq. (14) rapidly decays to zero with distance R for reasonable values of ζ_c. Thus, we have

$$\tilde{V}_{\ell_c}(R) \approx \frac{2\pi\Gamma\left(\ell_c+\frac{3}{2}\right)}{(2\ell_c+1)R^{\ell+1}\zeta_c^{\ell_c+\frac{3}{2}}} \tag{15}$$

Inclusion of the normalization factor from Eq. (7) gives

$$V_{\ell_c}(R) = \tilde{V}_{\ell_c}(R)N_{\ell_c 0}(\zeta_c) \tag{16}$$

$$\approx \left(\frac{2}{\zeta_c}\right)^{\frac{2\ell_c+3}{4}}\sqrt{\frac{2\pi\Gamma\left(\ell_c+\frac{3}{2}\right)}{2\ell_c+1}}\ R^{-\ell_c-1} \tag{17}$$

$$= \frac{(2\pi)^{\frac{3}{4}}\sqrt{(2\ell_c-1)!!}}{\zeta_c^{\frac{2\ell_c+3}{4}}R^{\ell_c+1}} \tag{18}$$

To arrive at our long-range estimate for $(ab|c)$, we simply scale Eq. (18) by the magnitude of the charge term from the bra to get

$$|(ab|c)| \approx |S_{ab}|\, V_{\ell_c}(R), \qquad \text{(large } R\text{)} \tag{19}$$

$$\approx |S_{ab}|\,\frac{(2\pi)^{\frac{3}{4}}\sqrt{(2\ell_c-1)!!}}{\zeta_c^{\frac{2\ell_c+3}{4}}R^{\ell_c+1}}, \quad (\ell_a=\ell_b=0). \tag{20}$$

We will refer to this formula as the SVℓ estimator. We derived it by considering the classical limit of the interaction of the electrostatic potential of the $|X)$ ket with the point-charge representation of the $(ss|$ ket; however, since the leading-order multipole of a product of two arbitrary Gaussians of any angular momenta is also a charge, at a large-enough distance away this estimate should also be sufficiently accurate.

C. Extension to Arbitrary ℓ_a and ℓ_b

Ideally, we would like to replace the charge-like contribution from the bra in Eq. (20) with a multipole expansion of the bra charge distribution to handle the general case. Unfortunately, it is too expensive to expand the bra distribution to multipoles of high enough order simply for the purpose of estimating integrals. Almost a decade of research on this topic[9,12,13,17,35–37] has led Ochsenfeld and co-workers to the conclusion that the best substitute for the multipole expansion in this context is the Schwarz bound,

$$Q_{ab} = |(ab|ab)|^{\frac{1}{2}} \tag{21}$$

(and hence the QQR estimator in Eq. (3)).[12] In the $(ss|$ case, the Schwarz estimate can be directly related to the overlap by

$$S_{ab} = Q_{ab}\left(\frac{\pi}{2(\zeta_a+\zeta_b)}\right)^{\frac{1}{4}}. \tag{22}$$

For higher angular momenta ℓ_a and ℓ_b, though, Q_{ab} incorporates contributions from higher-order multipoles, thus improving the estimate (as discussed in ref. 12). Incorporating this prefactor, we arrive at the estimator

$$|(ab|c)| \approx Q_{ab}\,\frac{\pi\sqrt{2(2\ell_c-1)!!}}{R^{\ell_c+1}\zeta_c^{\frac{2\ell_c+3}{4}}(\zeta_a+\zeta_b)^{\frac{1}{4}}}, \quad \text{(large } R\text{)} \tag{23}$$

which we call the QVℓ estimator. Again, we wish to estimate entire shells for practical purposes, so we have taken shell-wise Frobenius norms to obtain a Q_{ab} that is rotationally invariant:

$$Q_{ab} = \sqrt{\sum_{m_a=-\ell_a}^{\ell_a}\sum_{m_b=-\ell_b}^{\ell_b}(Q_{a(m_a),b(m_b)})^2}, \tag{24}$$

where $a(m_a)$ and $b(m_b)$ are the functions with L_z quantum numbers m_a and m_b in the *shells* with indices a and b. Note that this enforces rotational invariance since the Frobenius norm is invariant under rotations. It can easily be shown that when $\vec{R}_a \neq \vec{R}_b$, the GPT product distribution will always have an $\ell = 0$ contribution. Thus, for large enough R, the charge-like term of the bra multipole expansion will dominate, and the integral will eventually decay as $R^{-\ell_c+1}$, since the denominators of higher-order multipole terms will become much larger than this term.

D. Defining "Large R"

Up to this point, we have been vague about the definition of "large R," stating only that it is a distance from which a bra of the form $(ss|$ may be approximated as a point charge to sufficient precision. To a first approx-

imation, this concept is already well established in the context of CFMM.[14] In CFMM, two distributions Ω_{ab} and Ω_{cd} are considered "well-separated" if

$$R_{ab,cd} > \mathrm{ext}_{ab} + \mathrm{ext}_{cd} \tag{25}$$

where the extents ext_{ab} and ext_{cd} are given by

$$\mathrm{ext}_{ab} = \sqrt{\frac{2}{\zeta_a + \zeta_b}}\ \mathrm{erfc}^{-1}(\vartheta_{\mathrm{ws}}) \tag{26}$$

for some given WS threshold ϑ_{ws}. While this formula is obtained for spherical Gaussians, this definition turns out to be sufficient for our purposes, given the other approximations involved in our QVℓ estimator. For distributions that are not well separated, the QVℓ estimator reverts to Schwarz screening (Eq. (2)). Thus, the QVℓ estimate can be summarized as

$$|(ab|c)| \approx \begin{cases} Q_{ab}\,\dfrac{\pi\sqrt{2(2\ell_c-1)!!}}{R^{\ell_c+1}\zeta_c^{\frac{2\ell_c+3}{4}}(\zeta_a+\zeta_b)^{\frac{1}{4}}}, & R > \mathrm{ext}_{ab} + \mathrm{ext}_c \\ Q_{ab}Q_c, & R \le \mathrm{ext}_{ab} + \mathrm{ext}_c \end{cases} \tag{27}$$

E. Combining S_{ab} and Q_{ab}

As demonstrated in Section IV, the QVℓ estimator is robust, yielding estimates within a factor of 10 or so of the exact value for the vast majority of three-center ERIs. However, there are several aspects of the formulation that are fundamentally dissatisfying. While Q_{ab} has replaced S_{ab} to better account for the effects of higher order multipoles, the distance scaling factor remains that of the zeroth order multipole—namely, the overlap. But the reasoning behind this was that the term with slowest decay should dominate for large enough R. Hence S_{ab} should be a better representation of the bra contribution than Q_{ab} at large-enough separations. The problem is that when a and b are close together but differ significantly in angular momentum, the threshold beyond which higher order multipoles are negligible is much larger than $\mathrm{ext}_{ab} + \mathrm{ext}_c$. The CFMM extents only indicate when it is safe to approximate each term in the multipole expansion of the integral by point multipole interactions, not where it is necessarily safe to truncate the multipole expansion at the leading-order term. When the separation is such that higher order multipoles *are* actually small enough compared to the overlap, the SVℓ estimator is much better than the QVℓ estimator, since it gives the proper prefactor to the proper term. Incorporating this concept into the extents for the purposes of thresholding would require the computation of higher order multipole integrals, which we have already noted is too expensive for our purposes. Instead, one can roughly determine the importance of higher order multipoles by taking the ratio S_{ab}/Q_{ab}. When this ratio is small, higher-order multipole effects will overcome the additional R factors in the denominators of the hypothetical multipole expansion, warning us that the exclusion of the higher-order contributions from Q_{ab} could be dangerous. However, if this ratio is large enough, the dominant contribution to the multipole expansion will be the overlap, and S_{ab} should be used to approximate the bra contribution. We conclude that the estimator should be controlled by an additional screening parameter, ϑ_{SQ}. Defining a common prefactor for notational convenience:

$$\beta_\ell(\zeta) \equiv \zeta^{-\frac{2\ell+3}{4}}\sqrt{(2\ell-1)!!}, \tag{28}$$

we can now introduce our best estimator for three-center ERIs in terms of performance and flexibility, which we will call the SQVℓ estimator:

$$|(ab|c)| \approx \begin{cases} |S_{ab}|\dfrac{(2\pi)^{3/4}\,\beta_{\ell_c}(\zeta_c)}{R^{\ell_c+1}} & R > \mathrm{ext}_{ab} + \mathrm{ext}_c \\ & \text{and } S_{ab}/Q_{ab} > \vartheta_{\mathrm{SQ}} \\ Q_{ab}\dfrac{\pi\sqrt{2}\,\beta_{\ell_c}(\zeta_c)}{(\zeta_a+\zeta_b)^{\frac{1}{4}}R^{\ell_c+1}} & R > \mathrm{ext}_{ab} + \mathrm{ext}_c \\ & \text{and } S_{ab}/Q_{ab} \le \vartheta_{\mathrm{SQ}} \\ Q_{ab}Q_c & R \le \mathrm{ext}_{ab} + \mathrm{ext}_c\,. \end{cases} \tag{29}$$

In other words, the SQVℓ estimator "interpolates" between the SVℓ and QVℓ estimators: in the limits $\vartheta_{\mathrm{SQ}} \to 0$ and $\vartheta_{\mathrm{SQ}} \to \infty$ the SQVℓ estimator becomes equivalent to the SVℓ and QVℓ estimators, respectively.

As an aside, we note that for the purposes of our discussion here, three-center ERIs do not include cases where $\vec{R}_\mu = \vec{R}_\nu$ coincidentally. While these integrals are indeed part of the full three-center ERI tensor, they are a small enough part that the discussion of these integrals can be neglected in the current context. The SQVℓ estimator gives an approximate bound for these integrals, but because of the angular momentum addition rules, the actual decay with distance is sometimes much more rapid than the estimate accounts for, leading to significant overestimation. A better estimate for the special two-center case could be developed, but from a practical standpoint it is not worth the effort.

F. Contracted Basis Functions

Thus far, our discussion has focused only on primitive basis functions. For practical purposes, the extension to contracted basis functions is unimportant for many basis sets, since in many cases contracted basis functions are used to represent core orbitals, which do not contribute significantly to long-range integrals. Nevertheless, the extension of the SQVℓ estimator to contracted basis functions is relatively trivial, and yet it performs reasonably well even for basis sets composed entirely of contracted functions (see Section IV). Following Ochsenfeld, et al.,[12]

we typically define contracted extents $\mathrm{ext}_{\mu\nu}$ as

$$\mathrm{ext}_{\mu\nu} = \max_{a\in\mu,b\in\nu} \{\mathrm{ext}_{ab} + r_{ab,\mu\nu}\}, \tag{30}$$

where $r_{ab,\mu\nu}$ is the distance from the GPT center of the primitive pair $|ab)$ and the coefficient weighted center of charge of the product $|\mu\nu)$. However, in the case of generally contracted basis sets such as the ano-pVXZ series,[38] this formulation will substantially overestimate most of the extents, and a more careful (but less safe) formula is needed. For these basis sets, we use a coefficient-weighted average of the primitive pair extents:

$$\mathrm{ext}_{\mu\nu} = \frac{\sum_{a\in\mu,b\in\nu} c_{a,\mu}c_{b,\nu}(\mathrm{ext}_{ab} + r_{ab,\mu\nu})}{\sum_{a\in\mu,b\in\nu} c_{a,\mu}c_{b,\nu}} \tag{31}$$

where $c_{a,\mu}$ and $c_{b,\nu}$ are contraction coefficients. For the purposes of determining ζ_a, ζ_b, and ζ_c for, e.g., Eq. (29), the most diffuse exponent in the contraction is used for both standard and generally contracted basis sets. It is possible that a more efficient "effective exponent" formula could be developed, but a thorough investigation of screening for generally contracted basis sets is beyond the scope of this work.

III. COMPUTATIONAL DETAILS

The SVℓ, QVℓ, and SQVℓ estimators were implemented in a development version of the Massively Parallel Quantum Chemistry (`MPQC`)[39] package. We tested our estimates on a test set of three different molecular systems—benzene tetramer (π-stacked geometry with an inter-monomer separation of 3.2Å), linear icosane, and a cluster of 29 water molecules (Cartesian coordinates in supplemental information)—with five different basis set/auxiliary basis set pairs: cc-pVDZ[40] with cc-pVTZ/JK,[41] cc-pVTZ with cc-pV5Z/JK, Def2-SVP[42] with Def2-SVP/C,[43,44] aug-cc-pVTZ with aug-cc-pVTZ-RI,[45] and ano-pVDZ with aug-ano-pVTZ.[38] Since the chemistry of the molecules tested is less relevant to the current context (since the density matrix or other chemically important quantities are not involved in these estimates), the variety of basis sets with a variety of coefficients, exponents, and contraction schemes is more important to assessing the quality of the estimators than the variety of molecules examined. All error statistics were assessed with respect to shell-wise Frobenius norms for both the estimates and the actual integral values. All bra shell pairs were prescreened with a Schwarz threshold of 10^{-10}; that is, a bra pair $(\mu\nu|$ was excluded from all statistics if

$$|(\mu\nu|\mu\nu)|^{\frac{1}{2}} < \frac{10^{-10}}{\max|(X|X)|^{\frac{1}{2}}}. \tag{32}$$

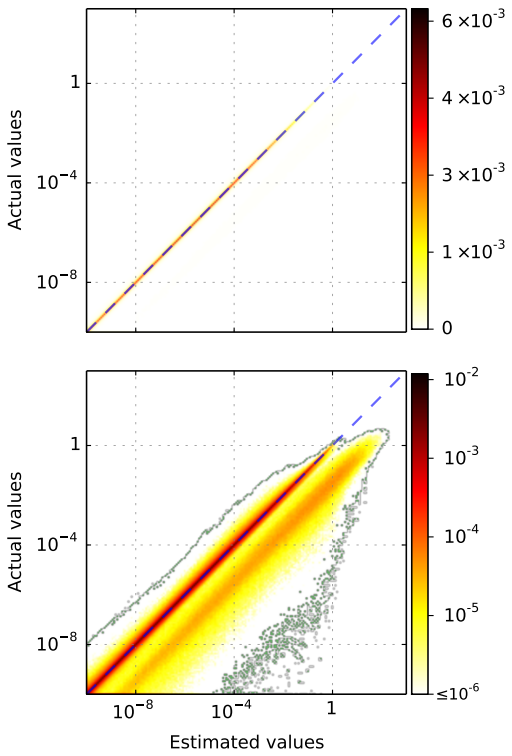

FIG. 1: Heat map of SQVℓ estimated versus exact integral shell norms obtained for our standard test set of molecules with the cc-pVTZ/cc-pV5Z/JK bases, with $\vartheta_{\mathrm{ws}} = 10^{-4}$ and $\vartheta_{\mathrm{SQ}} = 10^{-1}$. The plots use a linear (top) and logarithmic (bottom) color scale, with color values representing the fraction of total integrals in a given 2D histogram bin (of which there are 200 horizontal and 200 vertical). The green line in the bottom plot shows the boundary between the mostly empty histogram bins and the completely empty bins.

While the choice of this pair prescreening threshold has a small effect on the averages and standard deviations of the statistics, the worst case behaviors are largely unaffected by this choice, since these usually arise from pairs composed of functions with different angular momenta on neighboring atoms. These pairs usually have relatively large Schwarz estimates and are not affected by pair prescreening.

IV. RESULTS AND DISCUSSION

A. Qualitative Performance

Figure 1 shows a heat map of the estimated versus exact shell norms obtained with the SQVℓ estimator ($\vartheta_{\mathrm{ws}} = 10^{-4}$, $\vartheta_{\mathrm{SQ}} = 10^{-1}$). The data were obtained for our standard test set of molecules and the cc-pVTZ/cc-pV5Z/JK basis set pair. Using a linear color scale (top

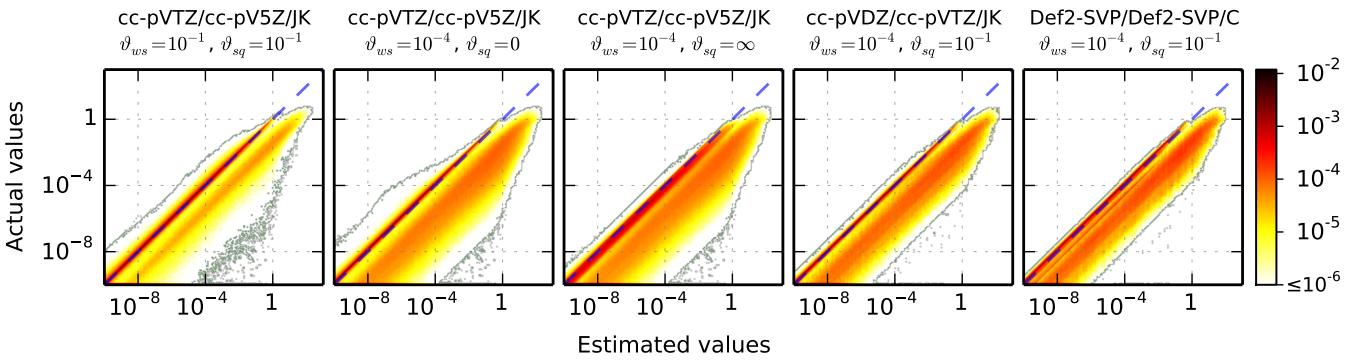

FIG. 2: The same plot from Figure 1 but with a different $\vartheta_{\rm ws}$, $\vartheta_{\rm SQ}$, or basis set. See caption of Figure 1 for details, and note that the color scale here is logarithmic. Note also the slightly expanded color scale relative to Figure 1 needed to accommodate the greater variety of basis sets.

plot), the estimator appears perfect; i.e., all data points fall along the ideal $I_{\rm estimate} = I_{\rm actual}$ line. A more complete picture of performance appears only with a logarithmic color scale (bottom plot). There are two bands in the plot. The first band follows closely the ideal line and contains the vast majority of the data points; it can be identified with the well-separated integrals (cases 1 and 2 in Eq. (29)). The minor band below the ideal line contains the data points where the estimated values are significantly greater than the exact ones, which is typical of the Schwarz-estimated integrals (case 3 in Eq. (29)). Note that the SQVℓ estimator eliminates the vast majority of the overestimates that would occur with the pure Schwarz screening. With still larger molecules or less diffuse basis sets, the Schwarz band would be even less prominent.

Since SQVℓ is not an upper bound, a few integrals are *under*estimated (the data above the ideal line); however this is an exceedingly rare occurrence. Keep in mind that the cc-pVTZ principal basis set and, particularly, the cc-pV5Z/JK auxiliary basis set are much larger and utilize much higher angular momentum than those typically used for large molecule computations. Indeed, the underestimated data points are almost not visible on heat maps for cc-pVDZ/cc-pVTZ/JK and Def2-SVP/Def2-SVP/C (fourth and fifth plots, Figure 2); i.e., these plots look similar that for a hypothetical rigorous upper bound.

The second and third plots in Figure 2 show the difference between the pure SVℓ estimator ($\vartheta_{\rm SQ} = 0$) and the pure QVℓ estimator ($\vartheta_{\rm SQ} = \infty$). The QV$\ell$ shows much less density of integrals in the underestimation region above the main diagonal, but at the expense of a noticeable broadening of the estimates along the main diagonal relative to the SVℓ or SQVℓ estimators. Finally, we note that in the first plot of Figure 2, the increase in the WS parameter relative to the data in Figure 1 does not change the picture much at all. For the molecules in our test set, the majority of the integrals that are well-separated with $\vartheta_{\rm ws} = 10^{-1}$ are also well-separated with $\vartheta_{\rm ws} = 10^{-4}$, so the difference is indistinguishable in this representation of the data. However, the outer limits indicating the worst over- and underestimates are basically identical for both $\vartheta_{\rm ws}$ values; this suggests that the worst case behavior can be attributed to a poor representation of higher-order multipoles in the bra rather than a lack of well-separatedness.

Figure 3 shows the data from Figure 1 split by angular momentum of the ket $|X)$. Notice that the main diagonal broadens with increasing ℓ_X. This is attributed to higher-order multipoles from the bra becoming more relevant with respect to their interaction with the ket as ℓ_X increases. In other words, integrals with larger ℓ_X are inherently harder to estimate, which is not surprising given that they are *significantly* more expensive to compute. As expected, the Schwarz band broadens and shifts further from the ideal as ℓ_X increases, because the distance factor (omitted from the Schwarz estimate) becomes more important to an accurate estimate with increasing ℓ_X.

B. Quantitative Performance

Performance of the integral estimators can be measured by analyzing large samples of the ratio of the estimate to the actual integral value:

$$F = I_{\rm estimate}/I_{\rm actual}, \tag{33}$$

where $I_{\rm estimate}$ and $I_{\rm actual}$ are the norms for the estimator and the actual computed integral shells, respectively (Ref. 12 used symbol F to denote this ratio, hence we will follow this notation for consistency).

Table 1 shows several statistical measures of F obtained from all unique shell triplets generated from our molecular test set with several different basis sets, and several relevant values of the estimator parameters ($\vartheta_{\rm ws}$ and $\vartheta_{\rm SQ}$). First, the estimator usually performs worse for larger basis sets than the smaller ones, as noted earlier. The exception is the cc-pVDZ/cc-pVTZ/JK pair, which outperforms the Def2-SVP/Def2-SVP/C pair in terms of $\bar{F}$, F_{max}, and $\sigma(\log F)$ (though the F_{min} values are slightly better for the latter pair). This excep-

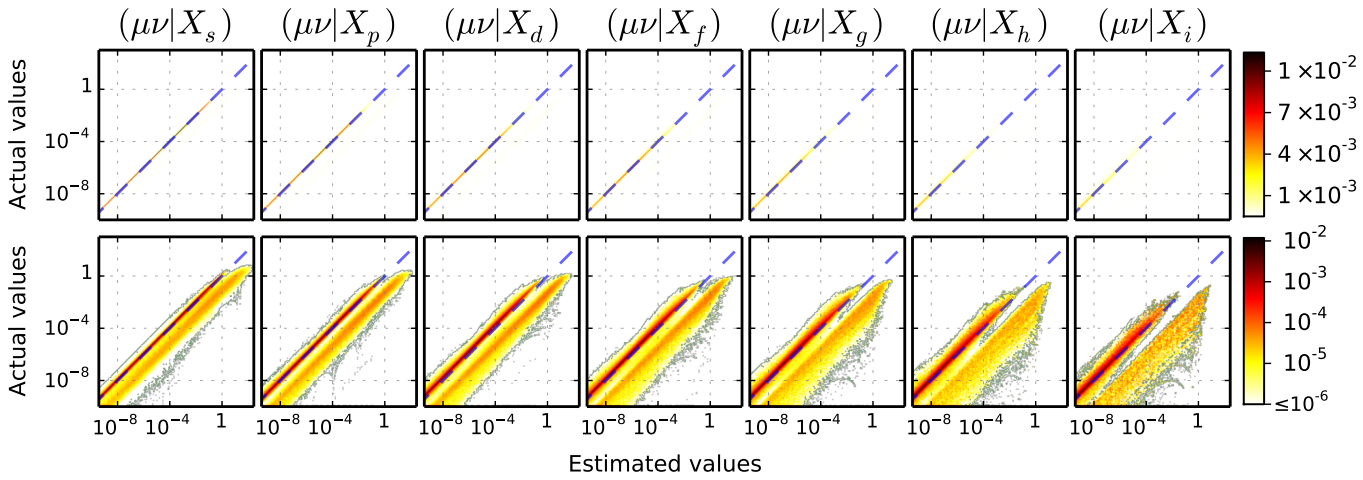

FIG. 3: Data from Figure 1 split by the angular momentum ℓ_X of the auxiliary basis function $|X)$. The color shows the fraction of integrals *with a given* ℓ_X in a particular histogram bin (rather than the fraction of all integrals, as in Figure 1). See caption of Figure 1 for details.

tion is attributed to the presence of contracted functions in the Def2-SVP/C basis, while the cc-pVTZ/JK basis is completely uncontracted. Our simple handling of contracted estimates (as discussed in Section II F) is to blame here. One could use the estimator on individual primitive triplets and then carry out the contraction, but we feel that this is a significant increase in effort for only a marginal increase in performance. Indeed, though the ano-pVDZ/aug-ano-pVTZ results are by far the worst, they are impressively tenable given the massive simplification from Section II F (massive in the context of generally contracted basis sets, that is).

Also note that while the average and standard deviation improve noticeably with the reduction of the ϑ_{SQ} parameter, the F_{min} gets substantially worse. This is a result of the exclusion of higher order effects from the QVℓ estimator which compensate for poor behavior in the worst edge cases at the expense of the average case. Similarly, the decrease in the ϑ_{ws} parameter comes at the cost of a roughly 30-40% decrease in the number of well-separated integrals (and thus, in the number of integrals accessible to the distance-dependent part of the estimator).

Table 2 shows F statistics for the shell triplets of cc-pVTZ/cc-pV5Z/JK grouped according to the ket angular momentum ℓ_X. Due to the narrower variety of exponents and contraction schemes for the higher angular momentum in the cc-pV5Z/JK basis set, higher angular momentum estimates are more statistically accurate with respect to mean and F_{max}, though the former is likely also a result of some error cancellation between over- and underestimates (as evidenced by the $\vartheta_{\text{SQ}} = 0$ and $\vartheta_{\text{SQ}} = 0.1$ cases). However, the minimum ratios for higher angular momentum are lower. This latter behavior is anticipated, since one would expect the importance of higher order multipoles in the bra to be more pronounced for integrals with higher angular momentum in the ket. Also, as the parameter ϑ_{SQ} increases, so does the consistency of the standard deviation with respect to angular momentum. For the pure SVℓ case, we see a much broader distribution for $\ell_X = 6$ than for $\ell_X = 0$, while in the pure QVℓ case, the standard deviation is nearly identical for all angular momenta. Again, this is attributed to the greater importance of higher-order effects for larger angular momenta.

C. Performance Versus Bra-Ket Distance

Figure 4 shows a 2D histogram of the ratio F plotted against bra–ket separation R_{AB}, with a linear color scale in the top plot and a logarithmic color scale on the bottom plot (note also the difference in the vertical scales). Figure 5 shows the same data split across ket angular momentum ℓ_X. In both figures, each distance bin is normalized individually to the range $[0, 1]$. As with Figure 1, the linear color scale is relatively uninformative, except insofar as it shows that the vast majority of well-separated integrals are estimated almost exactly at all distances. The logarithmic color scale reveals that most of the overestimates occur for shorter separations and taper off with increasing distance. The same trend can particularly be seen for higher angular momentum in Figure 5. Again, this shows that the higher angular momentum contributions to the bra are more important for integrals with higher ket angular momentum. At large bra-ket distances the estimator becomes more accurate due to the faster asymptotic decay of the contribution to the integral due to the higher angular momentum components to the bra.

TABLE 1: Statistics for the ratio $F = I_{\text{estimate}}/I_{\text{actual}}$ for our test set of three molecules (see text) and various basis sets, ϑ_{ws} values, and ϑ_{SQ} values. Note that the $\vartheta_{\text{SQ}} = 0$ case corresponds to the pure SVℓ estimator (see Section II B) and the $\vartheta_{\text{SQ}} = \infty$ case corresponds to the pure QVℓ estimator (see Section II C). Note that, while N_{ws} should remain exactly constant for a given ϑ_{ws}, in practice it varies slightly, because cases where the Schwarz estimate is smaller than the SQVℓ estimate are not counted as well-separated.

Basis	Aux. Basis	ϑ_{ws}	ϑ_{SQ}	$\bar{F}$	$\sigma(\log F)$	F_{min}	F_{max}	$N_{\text{ws}}/10^6$
cc-pVDZ	cc-pVTZ/JKFIT	10^{-1}	0	1.004	0.078	0.036	6.726	44.1
		10^{-1}	0.1	1.022	0.088	0.036	24.515	44.1
		10^{-1}	0.5	1.303	0.175	0.044	24.515	44.3
		10^{-1}	∞	1.542	0.166	0.050	24.515	44.6
		10^{-4}	0	1.001	0.058	0.083	3.080	31.0
		10^{-4}	0.1	1.023	0.074	0.199	24.485	31.0
		10^{-4}	0.5	1.336	0.175	0.199	24.485	31.2
		10^{-4}	∞	1.576	0.162	0.295	24.485	31.4
cc-pVTZ	cc-pV5Z/JKFIT	10^{-1}	0	0.995	0.117	2.8×10^{-4}	2293.872	177.6
		10^{-1}	0.1	1.143	0.142	0.002	3.7×10^4	177.6
		10^{-1}	0.5	1.907	0.249	0.002	3.7×10^4	180.4
		10^{-1}	∞	2.081	0.222	0.003	3.7×10^4	181.2
		10^{-4}	0	0.990	0.087	0.001	2293.872	129.9
		10^{-4}	0.1	1.175	0.130	0.037	3.7×10^4	130.0
		10^{-4}	0.5	2.004	0.247	0.042	3.7×10^4	132.4
		10^{-4}	∞	2.167	0.218	0.042	3.7×10^4	133.0
Def2-SVP	Def2-SVP/C	10^{-1}	0	1.459	0.211	0.103	20.615	28.8
		10^{-1}	0.1	1.483	0.215	0.144	128.377	28.8
		10^{-1}	0.5	1.927	0.264	0.198	128.377	28.9
		10^{-1}	∞	2.196	0.259	0.198	128.377	29.1
		10^{-4}	0	1.446	0.203	0.138	12.000	19.3
		10^{-4}	0.1	1.478	0.208	0.390	128.377	19.3
		10^{-4}	0.5	1.970	0.262	0.456	128.377	19.4
		10^{-4}	∞	2.240	0.255	0.514	128.377	19.5
aug-cc-pVTZ	aug-cc-pVTZ-RI	10^{-1}	0	1.002	0.159	7.4×10^{-5}	2293.872	532.1
		10^{-1}	0.1	1.846	0.245	0.011	4.7×10^4	532.8
		10^{-1}	0.5	2.726	0.311	0.012	4.7×10^4	538.4
		10^{-1}	∞	2.986	0.271	0.018	4.7×10^4	540.5
		10^{-4}	0	0.986	0.128	1.8×10^{-4}	2293.872	336.8
		10^{-4}	0.1	2.130	0.260	0.038	4.7×10^4	336.8
		10^{-4}	0.5	3.129	0.324	0.111	4.7×10^4	341.2
		10^{-4}	∞	3.368	0.282	0.122	4.7×10^4	342.6
ano-pVDZ	aug-ano-pVTZ	10^{-1}	0	3.770	0.274	0.001	2.4×10^5	65.7
		10^{-1}	0.1	3.866	0.268	0.024	2.4×10^5	65.7
		10^{-1}	0.5	4.659	0.292	0.036	3.5×10^5	66.6
		10^{-1}	∞	5.784	0.304	0.060	6.1×10^5	66.4
		10^{-4}	0	3.639	0.267	0.001	6.9×10^4	49.6
		10^{-4}	0.1	3.739	0.264	0.024	6.9×10^4	52.0
		10^{-4}	0.5	4.569	0.290	0.036	3.5×10^5	49.7
		10^{-4}	∞	5.696	0.298	0.060	3.5×10^5	49.5

V. CONCLUSIONS

We have introduced the SQVℓ estimator for three-center ERIs. It is exact for some classes of integrals and is confirmed numerically to provide very tight estimates of the integrals for a wide variety of basis set types. The estimator incorporates the correct leading-order dependence on the bra-ket distance, thus significantly increasing the sparsity of the three-center ERI tensor in reduced-scaling electronic structure methods. Computing the SQVℓ estimate for a given shell triplet is relatively cheap (on the order of a couple dozen CPU clock ticks; a detailed analysis is dependent on the implementation details). Thus many algorithms will see modest performance gains simply by incorporating the estimate directly into the integral computation without any index

TABLE 2: Statistics for the ratio $F = I_{\text{estimate}}/I_{\text{actual}}$ for our test set of three molecules with various ϑ_{SQ} parameters and split by angular momentum ℓ_X of the auxiliary ket shell $|X)$, with the cc-pVTZ[40] basis and the cc-pV5Z/JK[41] auxiliary basis and $\vartheta_{\text{ws}} = 10^{-4}$.

ϑ_{SQ}	ℓ_X	$\bar{F}$	$\sigma(\log F)$	F_{min}	F_{max}	$N_{\text{ws}}/10^6$
	0	1.009	0.047	0.013	2293.87	43.5
	1	0.996	0.067	0.006	19.20	27.8
	2	0.987	0.087	0.004	693.32	24.1
0	3	0.975	0.108	0.002	28.33	17.2
	4	0.962	0.130	0.002	16.88	10.9
	5	0.943	0.158	0.001	15.72	5.3
	6	0.920	0.184	0.001	9.25	1.2
	0	1.282	0.122	0.296	3.7×10^4	43.5
	1	1.176	0.122	0.145	345.38	27.8
	2	1.134	0.127	0.100	693.32	24.1
0.1	3	1.091	0.133	0.080	334.74	17.2
	4	1.057	0.144	0.062	312.09	10.9
	5	1.026	0.162	0.042	269.45	5.3
	6	0.992	0.182	0.037	165.29	1.2
	0	2.112	0.247	0.500	3.7×10^4	44.1
	1	2.006	0.246	0.282	345.38	28.2
	2	1.977	0.246	0.193	4299.39	24.6
0.5	3	1.917	0.246	0.111	334.74	17.5
	4	1.877	0.248	0.071	312.09	11.2
	5	1.842	0.251	0.042	269.45	5.5
	6	1.800	0.258	0.064	165.29	1.2
	0	2.275	0.218	0.500	3.7×10^4	44.3
	1	2.169	0.217	0.282	345.38	28.4
	2	2.138	0.217	0.204	4299.39	24.7
∞	3	2.078	0.217	0.121	334.74	17.6
	4	2.038	0.218	0.071	312.09	11.3
	5	2.003	0.221	0.042	269.45	5.6
	6	1.958	0.227	0.064	165.29	1.3

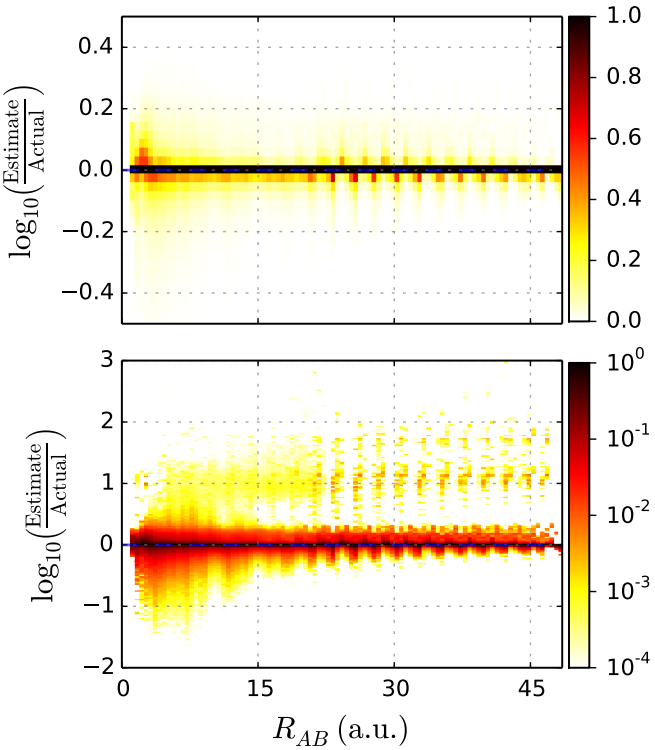

FIG. 4: Heat map of the ratio $F = I_{\text{estimate}}/I_{\text{actual}}$ (well-separated integrals only) for our test set of molecules using the cc-pVTZ/cc-pV5Z/JK basis, with $\vartheta_{\text{ws}} = 10^{-4}$ and $\vartheta_{\text{SQ}} = 10^{-1}$ using a linear (top) and logarithmic (bottom) color scale. Ratios are normalized within each of 100 distance histogram bins so that the maximum in any given column is 1.0. Note the difference in vertical scales between the two plots.

reordering or extra bookkeeping. More substantial gains are anticipated in algorithms that do not visit every shell triplet, even for the purposes of estimation.

Our basic tests performed outside of the context of any particular electronic structure method show that the norm of the most integrals in which the bra and ket are "well-separated" are estimated nearly perfectly, particularly for smaller basis sets that are more likely to be used for large molecule computations. Our tests also show that the extent by which the integral norms are underestimated can be reduced readily without incurring explosion in computational cost. Tuning the adjustable parameters of the estimator (ϑ_{ws} and ϑ_{SQ}) in the context of a given electronic structure method will be a function of the target precision and should be determined, as always in atomic basis electronic structure, by benchmarking. The use of the estimator for reduced-scaling construction of the Hartree-Fock exchange with the concentric atomic density fitting approximation[23] will be described in an upcoming manuscript.

VI. ACKNOWLEDGEMENTS

The work by DSH and EFV was supported by NSF grants CHE-1362655 and ACI-1047696, and Camille Dreyfus Teacher-Scholar Award. The work by DSH and HFS was supported by NSF grant CHE-1361178. This work used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U. S. Department of Energy under Contract No. DE-AC02-05CH11231.

[1] D. R. Hartree, Rep. Prog. Phys. **11**, 113 (1947).
[2] A. D. Becke, J. Chem. Phys. **98**, 1372 (1993).
[3] C. Møller and M. S. Plesset, Phys. Rev. **46**, 618 (1934).
[4] V. Dyczmons, Theor. Chim. Acta **28**, 307 (1973).
[5] M. Häser and R. Ahlrichs, J. Comput. Chem. **10**, 104 (1989).
[6] P. M. W. Gill, B. G. Johnson, and J. A. Pople, Chem. Phys. Lett. **217**, 65 (1994).
[7] D. L. Strout and G. E. Scuseria, J. Chem. Phys. **102**, 8448 (1995).
[8] M. Schütz, R. Lindh, and H.-J. Werner, Mol. Phys. **96**, 719 (1999).
[9] D. S. Lambrecht and C. Ochsenfeld, J. Chem. Phys. **123**, 184101 (2005).
[10] D. S. Lambrecht, B. Doser, and C. Ochsenfeld, J. Chem. Phys. **123**, 184102 (2005).

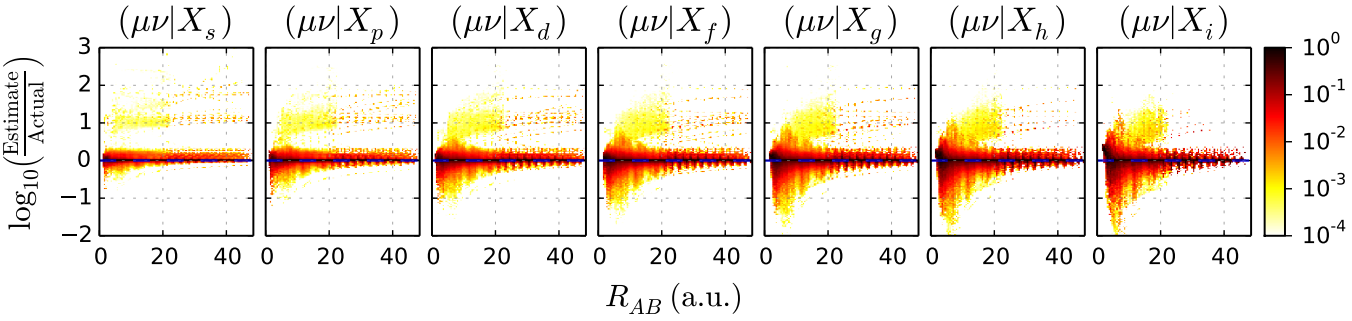

FIG. 5: Data from Figure 4 split by angular momentum of the auxiliary basis function $|X)$ using a logarithmic color scale. See caption of Figure 4 for details.

[11]B. Doser, D. S. Lambrecht, and C. Ochsenfeld, Phys. Chem. Chem. Phys. **10**, 3335 (2008).
[12]S. A. Maurer, D. S. Lambrecht, D. Flaig, and C. Ochsenfeld, J. Chem. Phys. **136**, 144107 (2012).
[13]S. A. Maurer, D. S. Lambrecht, J. Kussmann, and C. Ochsenfeld, J. Chem. Phys. **138**, 014101 (2013).
[14]C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, Chem. Phys. Lett. **230**, 8 (1994).
[15]E. H. Rubensson, E. Rudberg, and P. Sałek, J. Math. Phys. **49**, 032103 (2008).
[16]E. Rudberg, E. H. Rubensson, and P. Sałek, J. Chem. Theory Comput. **5**, 80 (2009).
[17]J. Kussmann, M. Beer, and C. Ochsenfeld, WIREs Comput. Mol. Sci. **3**, 614 (2013).
[18]D. R. Bowler and T. Miyazaki, Rep. Prog. Phys. **75**, 036503 (2012).
[19]C. Riplinger and F. Neese, J. Chem. Phys. **138**, 034106 (2013).
[20]A. Hansen, D. G. Liakos, and F. Neese, J. Chem. Phys. **135**, 214102 (2011).
[21]F. Neese, A. Hansen, and D. G. Liakos, J. Chem. Phys. **131**, 064103 (2009).
[22]F. Neese, F. Wennmohs, and A. Hansen, J. Chem. Phys. **130**, 114108 (2009).
[23]D. S. Hollman, J. J. Wilke, and H. F. Schaefer, J. Chem. Phys. **138**, 064107 (2013).
[24]A. Sodt, J. E. Subotnik, and M. Head-Gordon, J. Chem. Phys. **125**, 194109 (2006).
[25]T. B. Adler, H.-J. Werner, and F. R. Manby, J. Chem. Phys. **130**, 054106 (2009).
[26]T. B. Adler and H.-J. Werner, J. Chem. Phys. **130**, 241101 (2009).
[27]C. Hättig, D. P. Tew, and B. Helmich, J. Chem. Phys. **136**, 204105 (2012).
[28]D. P. Tew, B. Helmich, and C. Hättig, J. Chem. Phys. **135**, 074107 (2011).
[29]D. Usvyat, J. Chem. Phys. **139**, 194101 (2013).
[30]T. B. Adler and H.-J. Werner, J. Chem. Phys. **135**, 144117 (2011).
[31]A. F. Izmaylov and G. E. Scuseria, Phys. Chem. Chem. Phys. **10**, 3421 (2008).
[32]F. Aquilante, L. Gagliardi, T. B. Pedersen, and R. Lindh, J. Chem. Phys. **130**, 154107 (2009).
[33]T. Helgaker, P. Jørgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (Wiley, Chichester New York, 2000).
[34]J. F. Rico, R. López, I. Ema, and G. Ramírez, Int. J. Quantum Chem. **113**, 1544 (2013).
[35]C. Ochsenfeld and J. Kussmann, Rev. Comp. Chem. (2007).
[36]J. Kussmann and C. Ochsenfeld, J. Chem. Phys. **127**, 204103 (2007).
[37]M. Beer and C. Ochsenfeld, J. Chem. Phys. **128**, 221102 (2008).
[38]F. Neese and E. F. Valeev, J. Chem. Theory Comput. **7**, 33 (2011).
[39]C. Janssen, E. Seidl, and M. Colvin, in *ACS Symposium Series, Parallel Computing in Computational Chemistry*, Vol. 592 (1995) pp. 47–61.
[40]T. H. Dunning, J. Chem. Phys. **90**, 1007 (1989).
[41]F. Weigend, Phys. Chem. Chem. Phys. **4**, 4285 (2002).
[42]F. Weigend and R. Ahlrichs, Phys. Chem. Chem. Phys. **7**, 3297 (2005).
[43]F. Weigend, M. Häser, H. Patzelt, and R. Ahlrichs, Chem. Phys. Lett. **294**, 143 (1998).
[44]A. Hellweg, C. Hättig, S. Höfener, and W. Klopper, Theor. Chim. Acta **117**, 587 (2007).
[45]F. Weigend, A. Köhn, and C. Hättig, J. Chem. Phys. **116**, 3175 (2002).